

Catalogue no. 15-206-X — No. 024
ISSN 1710-5269
ISBN 978-1-100-12632-6

Research Paper

The Canadian Productivity Review

Productivity Spillovers from Competitive Reallocation: Evidence from Canadian Manufacturing Plants

by Guy Gellatly and John Baldwin

Economic Analysis Division
R.H. Coats Bldg., 18th floor, 100 Tunney's Pasture Driveway
Ottawa, Ontario K1A 0T6

Telephone: 1-800-263-1136



 Statistics Canada Statistique Canada

Canada

How to obtain more information

Specific inquiries about this product and related statistics or services should be directed to the Media Hotline, Communications and Library Services Division, Statistics Canada, Ottawa, Ontario K1A 0T6 (telephone: 613-951-4636).

For information about this product or the wide range of services and data available from Statistics Canada, visit our website at www.statcan.gc.ca or contact us by e-mail at infostats@statcan.gc.ca or by telephone from 8:30 a.m. to 4:30 p.m. Monday to Friday:

Statistics Canada National Contact Centre

Toll-free telephone (Canada and the United States):

Inquiries line	1-800-263-1136
National telecommunications device for the hearing impaired	1-800-363-7629
Fax line	1-877-287-4369

Local or international calls:

Inquiries line	1-613-951-8116
Fax line	1-613-951-0581

Depository services program

Inquiries line	1-800-635-7943
Fax line	1-800-565-7757

Information to access the product

This product, Catalogue no. 15-206-X, is available for free in electronic format. To obtain a single issue, visit our website at www.statcan.gc.ca and select "Publications."

Standards of service to the public

Statistics Canada is committed to serving its clients in a prompt, reliable and courteous manner. To this end, the Agency has developed standards of service which its employees observe in serving its clients. To obtain a copy of these service standards, please contact Statistics Canada toll free at 1-800-263-1136. The service standards are also published on www.statcan.gc.ca under "About us" > "Providing services to Canadians."

The Canadian Productivity Review

The Canadian Productivity Review is a series of applied studies that address issues involving the measurement, explanation, and improvement of productivity. Themes covered in the review include, but are not limited to, economic performance, capital formation, labour, prices, environment, trade, and efficiency at both national and provincial levels. The Review publishes empirical research, at different levels of aggregation, based on growth accounting, econometrics, index numbers and mathematical programming. The empirical research illustrates the application of theory and techniques to relevant public policy issues.

The primary distribution medium for this series is the Internet. These studies can be downloaded from the Internet at www.statcan.gc.ca for free.

All papers in **The Canadian Productivity Review** go through institutional review to ensure that they conform to Statistics Canada's mandate as a government statistical agency and adhere to generally accepted standards of good professional practice.

The papers in the series often include results derived from multivariate analysis or other statistical techniques. It should be recognized that the results of these analyses are subject to uncertainty in the reported estimates.

The level of uncertainty will depend on several factors: the nature of the functional form used in the multivariate analysis; the type of econometric technique employed; the appropriateness of the statistical assumptions embedded in the model or technique; the comprehensiveness of the variables included in the analysis; and the accuracy of the data that are utilized. The peer group review process is meant to ensure that the papers in the series have followed accepted standards to minimize problems in each of these areas.

Productivity Spillovers from Competitive Reallocation: Evidence from Canadian Manufacturing Plants

Guy Gellatly and John Baldwin

Published by authority of the Minister responsible for Statistics Canada

© Minister of Industry, 2009

All rights reserved. The content of this electronic publication may be reproduced, in whole or in part, and by any means, without further permission from Statistics Canada, subject to the following conditions: that it be done solely for the purposes of private study, research, criticism, review or newspaper summary, and/or for non-commercial purposes; and that Statistics Canada be fully acknowledged as follows: Source (or “Adapted from,” if appropriate): Statistics Canada, year of publication, name of product, catalogue number, volume and issue numbers, reference period and page(s). Otherwise, no part of this publication may be reproduced, stored in a retrieval system or transmitted in any form, by any means—electronic, mechanical or photocopy—or for any purposes without prior written permission of Licensing Services, Client Services Division, Statistics Canada, Ottawa, Ontario, Canada K1A 0T6.

May 2009

Catalogue no. 15-206-X, no. 024
Frequency: Occasional

ISSN 1710-5269
ISBN 978-1-100-12632-6

Cette publication est disponible en français (n° 15-206-X au catalogue, n° 024).

Note of appreciation

Canada owes the success of its statistical system to a long-standing partnership between Statistics Canada, the citizens of Canada, its businesses, governments and other institutions. Accurate and timely statistical information could not be produced without their continued cooperation and goodwill.

Acknowledgements

The authors are indebted to Bob Gibson of the Economic Analysis Division for the preparation of data and to Desmond Beckstead and Mark Brown of the same division for their many helpful comments and suggestions. We are also grateful for comments that we received on earlier versions of this paper from participants at the Annual Meetings of the Canadian Economics Association (May 2005), at the Industry Canada Economic Seminar Series (January 2006), and at the North American Regional Science Council Annual Meetings (November 2008).

Table of contents

Abstract.....	6
Executive summary.....	7
1 Introduction.....	10
2 Productivity spillovers and competitive reallocation.....	11
3 Data and measurement.....	15
4 A model of productivity spillovers.....	19
5 Results and analysis.....	24
6 Conclusions.....	29
Appendix A Editing data samples for the regression analysis.....	32
References.....	34

Symbols

The following standard symbols are used in Statistics Canada publications:

- . not available for any reference period
- .. not available for a specific reference period
- ... not applicable
- 0 true zero or a value rounded to zero
- 0^s value rounded to 0 (zero) where there is a meaningful distinction between true zero and the value that was rounded
- ^p preliminary
- ^r revised
- x suppressed to meet the confidentiality requirements of the *Statistics Act*
- ^E use with caution
- F too unreliable to be published

Abstract

This paper uses plant-level data on productivity growth and changes in market share over different periods during the 1970s, 1980s, and 1990s to investigate whether plants with declining market shares obtain productivity spillovers from more successful producers and whether the impact of spillovers is affected by the distance between plants. We are primarily interested in the extent to which productivity externalities moderate the centrifugal forces that separate growing plants from declining rivals because of the productivity advantages enjoyed by the former.

The paper focuses on the productivity performance of plants with declining market shares—as potential receivers of productivity spillovers. Two possible sources for these spillovers are examined—rival plants operating at the technological frontier and rivals that are actively gaining market share. The analysis advances a model of the externality process in which the productivity of declining plants is influenced by (1) the economic distance of the declining plant from its technological frontier at the beginning of any period, (2) contemporaneous productivity gains in rival plants that are actively wresting market share away from decliners, and (3) the distance between rival plants.

We evaluate the existence and magnitude of these sources of spillovers—frontier plants and market-share gainers—because of what they reveal about the types of productive information that struggling plants may be able to assimilate from rivals. Spillovers from the plants at the existing frontier are likely to reflect the established best practices of industry leaders; spillovers coming from market-share gainers involve new sources of productive knowledge that emerge as the frontier is actively being re-established. Our model also incorporates geographic information on the proximity of declining plants to both frontier plants and market-share gainers to test whether productivity spillovers are spatially circumscribed. The results provide evidence that productivity improvements in more successful plants benefit their struggling rivals and that these benefits are inversely related to distance; however, the magnitude of spillovers from growing plants to decliners is relatively small. Spillovers do not offer much of a safety net for producers that are losing the productivity race. The paper also shows that declining plants that start out behind the technological frontier are likely to fall further behind, after the impact of mean reversion is taken into account.

Executive summary

Research on the dynamics of competition has shown that large amounts of market share are transferred over time from declining to growing firms. The firms that gain market share do so as the result of considerable gains in relative productivity. This reallocative process has been shown to contribute substantially to aggregate productivity growth.

The identity of growing firms and declining firms is difficult to predict *ex ante*. At the beginning of any time period, there is little difference in the productivity of those who will subsequently gain and those who will subsequently lose market share in the next period. But the former find ways to improve their productivity relative to the latter, and the sudden gain in relative productivity is translated into large gains in market share (Baldwin and Sabourin 2001, 2004; Baldwin and Gellatly 2003). The resulting shift of resources within an industry from less- to more-productive producers accounts for a considerable proportion of overall industry productivity growth—up to 50% in the manufacturing sector (Baldwin and Gu 2006) and around 100% in the retailing sector (Baldwin and Gu 2008).

Complementary research from Canadian business surveys has shown that this sudden emergence of differences in productivity is partly related to differences in the capabilities of firms to innovate and exploit the benefits of advanced technology (Baldwin and Sabourin 2001, 2004; Baldwin, Sabourin and Smith 2003). These emerging differences in productivity create centrifugal forces within the marketplace that increase the market share of those gaining relative productivity and decrease the market share for those with declining relative productivity. Opposing this tendency are centripetal forces within the marketplace that help those losing market share to attenuate their productivity differences relative to rivals that are surging ahead.

Externalities from spillovers contribute to these centripetal forces. Spillovers have long been regarded as an important feature of competitive markets. They involve the unpriced transfers of productive knowledge from some producers—firms that possess special assets, and/or organizational advantages—to rivals. The special knowledge of techniques, technologies, patents and organizational forms that permits some firms to develop an advantage can be transmitted via spillovers to those who fall behind. The magnitude and speed with which this transfer occurs will determine whether those who did not initially discover new techniques or new products survive in the market place.

The importance of these externalities will determine the speed with which populations turn over and renew themselves via entry and exit, and influence the extent to which existing entities can transform themselves in time to save themselves from having to exit. It will determine whether the ecological model of sorting is one where the less productive are rapidly culled or whether there is a type of safety net coming from spillovers that prevents them from falling too far behind the industry leaders. Finally, the speed with which the knowledge of growing firms is transferred to others may affect the incentive of high achievers to innovate and thus overall economic progress. The notion that externalities should be encouraged underlies many industrial policies that are aimed at facilitating the transfer of knowledge to those who are disadvantaged. This goes to the heart of patent laws that offer legal protection for ideas—but only for a prescribed period of time, so that the ideas behind patents can be disseminated as widely as possible.

Research from microeconomic panels has produced a considerable number of studies regarding the sources and types of centrifugal processes at work within the Canadian business sector. These are forces that encourage the competitive separation of more- from less-productive firms. Less effort has been spent examining the nature of the centripetal forces that may potentially mitigate the competitive separation of rival firms. This paper is our first attempt to look at this issue. It focuses on the relationship between the productivity growth of declining plants (possible receivers of spillovers) and that of growing plants (possible senders of spillovers). There are other externalities that are of potential interest, such as those that are transferred to successful plants, for example. But since so much attention has been placed on the process of competitive reallocation that shifts economic resources away from declining to growing firms, an examination of the productivity spillovers that link growers to decliners presents a logical starting point.

This paper uses plant-level data on productivity growth and changes in market share over three periods (1979 to 1984, 1984 to 1990 and 1990 to 1996) to investigate whether plants with declining market shares in each of these periods obtain productivity spillovers from rival producers that are **revealed by the market place** to be more successful. It examines the extent to which the productivity performance of plants with declining market shares is influenced by (1) their initial proximity to the existing technological frontier, and by (2) contemporaneous productivity changes in rival incumbents that are growing their market share. In doing so, it asks whether the impact of these spillovers is circumscribed by distance.

The paper evaluates these possible sources of spillovers—frontier plants at the beginning of a period and plants which gain market share over a period—because of what they reveal about the types of productive information that struggling plants may be able to assimilate from rivals. Spillovers from the existing frontier are likely to reflect the established best practices of industry leaders; those from market-share gainers involve new sources of productive knowledge that emerge as the frontier is actively being re-established. The spillover model advanced herein incorporates geographic information on the proximity of declining plants to both frontier plants and market-share gainers to test whether productivity spillovers are spatially circumscribed.

The analysis finds that plants losing market share do not obtain productivity spillovers from rivals located at the technological frontier. After mean reversion is taken into account, declining plants that start out further behind the frontier actually exhibit slower, not faster, rates of productivity growth—evidence that knowledge from the old frontier does not spill over to low-efficiency producers. The further a declining plant starts off behind the frontier, the lower its productivity growth is likely to be.

The analysis does find that plants with declining market shares obtain productivity spillovers from rivals gaining market share—plants with a revealed competitive advantage. These are rivals on the right side of the competitive struggle as the dynamics of competition are actively being played out. While declining plants obtain some measure of benefit from these rivals, productivity spillovers from market-share gainers to market-share decliners are spatially circumscribed and qualitatively modest.

At most, these productivity externalities from growing plants increase the rate of productivity growth of declining plants by between 0.18 and 0.34 of a percentage point. This is relatively small compared to the almost 4-percentage-point difference between the median growth rate of growing and declining plants, averaged over the three periods. The externalities measured here

then have a relatively minor impact on closing the gap between the more and the less successful. The evolutionary process resembles a contest that quickly rewards the more successful and leaves the less successful to fall further and further behind the leaders.

1 Introduction

Spillovers have long been regarded as an important feature of competitive markets. They involve the unpriced transfers of productive knowledge from some producers—firms that possess special assets, and/or organizational advantages—to rivals. The transmission of these benefits has been posited to occur via different mechanisms, ranging from demonstration effects in which rivals acquire and assimilate information on the best practices of industry leaders, to pro-competition effects in which the entry of highly productive firms induces rivals to improve their operating efficiency.¹ General descriptions of how and where this occurs are often static in nature—the recipients of spillovers are simply seen to obtain benefits from competitors because of the diffusion of new technologies or best practices. Such descriptions seldom make reference to the underlying forces of competitive restructuring that are continuously at work within industries—the ebb and flow of growth and decline as some producers benefit at the expense of others. As a consequence, little attempt has been made to either conceptualize productivity spillovers as part of the dynamic process that occurs as more successful producers wrest market share away from rivals or to quantify the impact of spillovers on this process.

This dynamic process is the essence of competition. Microeconomic research on the dynamics of the competitive process has shown that large amounts of market share are transferred over time from declining to growing firms. The firms that gain market share do so as the result of considerable gains in relative productivity. This occurs as these emerging differences in relative productivity between rival producers create centrifugal forces within industries that increase the market share of producers that are becoming more productive and that decrease the market share of those becoming less productive. This reallocative process, in turn, fosters substantial improvements in industrial productivity.

Spillovers are one factor that can be expected to affect the speed of this reallocative process. They may act as a centripetal force that counteracts the extent to which the competitive process separates growing firms from declining firms. At issue is whether these spillovers are large enough to contribute significantly to the productivity growth of plants in decline—whether the ecological model of sorting at work is one where the less productive are rapidly culled or whether there is a type of safety net coming from spillovers that prevents the laggards from falling too far behind the industry leaders.

Baldwin (1995), in his empirical analysis of competitive dynamics, raised the possibility that the transfer of market share from producers in decline to growing rivals involves these spillover effects—with some of the gains from better performers spilling back to laggards. But little evidence of these spillovers was apparent when the author examined correlations of average growth rates in these two groups across industries in the 1970s.

In what follows, we revisit the idea of productivity spillovers from competitive reallocation advanced in Baldwin (1995) using detailed plant-level regressions over a set of periods that span three decades. The analysis of competitive reallocation focuses on plants that are in decline—losing market share—and those that are growing—gaining market share. The former are most at

1. For a useful discussion of spillovers and possible transmission mechanisms, see Markusen's (1998) review of multinationals. See also Caves (1974), Globerman (1979) and Baldwin and Hanel (2003).

risk and, on balance, have much weaker productivity performance. These declining plants provide the focus of this paper, as they epitomize those that suffer the most from the re-allocative forces at work within the competitive system. Further, they potentially stand to gain the most from externalities since their productivity growth is generally far behind those that are gaining market share.

Using panel data on Canadian manufacturing plants, we test whether the productivity growth of plants in decline—incumbent producers that lose market share over different medium-run periods that span different decades—is influenced by (1) their initial proximity to the technological frontier, and by (2) contemporaneous productivity changes in rival incumbents that are growing their market share. We are interested in these two possible sources of spillovers—frontier plants and market-share gainers—because of what they reveal about the types of productive information that competitors may be able to acquire and assimilate from rivals. Older information from the frontier is more likely to reflect the established best practices of industry leaders and, as a consequence, may be less difficult for potential receivers of spillovers to obtain; by contrast, acquiring information from rivals that are actively capturing market share and growing their productivity reflects a type of dynamic learning that occurs as the frontier is continuously being re-established.

The paper is organized as follows. Section 2 describes the applied research on productivity that motivates the current analysis. Measurement issues are discussed in Section 3, followed by descriptive statistics on changes in market share and plant productivity. Our externality model is presented in Section 4. The results of our regression analysis are reported and discussed in Section 5. Section 6 concludes.

2 Productivity spillovers and competitive reallocation

Longitudinal research on productivity dynamics has found a strong link between competitive restructuring and productivity performance. In their survey article, Bartelsman and Doms (2000, p. 571) noted that these studies have aided in “resurrecting the Schumpeterian idea of creative-destruction,” as the process of competitive reallocation has been shown to represent “a major force contributing to (aggregate) productivity growth.” These studies portray competition as a dynamic process waged by heterogeneous rivals, one consistent with Schumpeterian notions of competitive replacement wherein more innovative, efficient producers supplant the less capable. One desirable feature of these evolutionary models of productivity growth is that individual producers are not reduced to the status of representative agents, consigned to some exogenously determined fate. Following the theoretical contributions on learning by Jovanovic (1982) and Ericson and Pakes (1995), new studies on productivity dynamics often portray firms as endogenous agents that learn and adapt in order to improve their relative productivity. This competitive turbulence among unequals serves, in part, to improve aggregate productivity performance.²

Research on the dynamics of competition in the Canadian manufacturing sector has shown that large amounts of market share are transferred over time from declining to growing firms. More

2. For an overview of micro-economic research on productivity dynamics, see Bartelsman and Doms (2000) that includes a more comprehensive discussion of many of these points. For a summary of the issues involved in measuring the impact of turnover on productivity growth, see Baldwin and Gu (2006).

importantly, this process underpins industry productivity growth. The firms that gain market share do so as the result of considerable gains in relative productivity.

The identification of growing firms and declining firms is difficult to predict *ex ante*. At the beginning of any time period, there is little difference in the productivity growth of those who will subsequently gain and those who will subsequently lose market share in the next period. But the former find ways to improve their productivity relative to the latter, and the sudden gain in relative productivity is translated into large gains in market share. This emergence of differences in productivity is partly related to differences in the capabilities of firms to innovate and adopt advanced technologies (Baldwin and Sabourin 2001, 2004; Baldwin, Sabourin and Smith 2003).

A recent study of Canadian manufacturers quantifies the contribution that competitive reallocation among rival incumbents makes to aggregate productivity growth in the 1980s and 1990s. Baldwin and Gu (2006) decompose productivity growth into three components: the amounts coming from, respectively, (1) organic productivity growth in market-share gainers, (2) organic productivity growth in market-share losers, and (3) the productivity growth that comes from transfers of market share from losers to gainers.³ The authors show that the transfer of market share to plants that are becoming more productive accounts for a large share of long-run productivity growth. They note that, from 1988 to 1997, some 53% of labour productivity growth in Canadian manufacturing was due to the competitive process that re-allocates output from less- to more-productive plants, with some 73% of this growth stemming from transfers among incumbents. Together these competitive transfers to market-share gainers, coupled with organic productivity gains in these more successful producers, amount to about 80% of total productivity growth.⁴

The competitive process creates productivity growth by rewarding more productive enterprises with increases in market share. But the decompositions that are used to estimate the contribution of competitive reallocation to productivity growth generally assume that productivity growth in different firms is independent. If those firms that are gaining market share are also creating spillovers for rivals losing market share, then part of the productivity growth in this latter group should also be attributed to market-share gainers, and the importance of this group would be even higher than the existing literature suggests. This is one of the issues that is addressed here.

The issue of spillovers is also germane to an important facet of applied research on productivity dynamics, namely, the need to better understand the interactions among rival producers and the nature of the ecological process that underpins growth dynamics. These interactions are, on balance, understudied. As Bartelsman and Doms (2000, p. 592) note, “the decisions and behavior of firms or establishments have been examined under the assumption of no interactions among firms,” despite the fact that “[t]he decision to use certain technologies or to increase output depends on the actions of competing firms.” Herein we focus on one type of rival-to-rival interaction—the possible spillovers that more successful growing producers confer on their less-successful declining rivals.⁵

3. Earlier studies (Baldwin 1995) find similar orders of magnitude.

4. The contribution of firm turnover to aggregate productivity growth is even higher in the retail sector (Baldwin and Gu 2008).

5. This characterization of spillovers as an involuntary transmission of benefits—productive knowledge not fully appropriated by the firm generating the externality—is found in Markusen (1998).

Some elaboration on the concept of spillovers examined herein is warranted. Spillovers entail the unpriced transfers of productive knowledge from certain firms—those that possess special assets, and/or organizational advantages—to rivals. A voluminous literature on spillovers has emerged, which focuses on measuring the technological benefits that foreign multinationals bring to host economies. Many of these studies investigate the effects of inbound foreign direct investment on domestic economies. Others focus on sectoral knowledge flows, concentrating on the non-excludability of intangible assets like research and development. Research on spillovers at the plant or establishment level from longitudinal data is less common, given the relative newness of many of these data sources. One example is Griffith, Redding and Simpson (2002), who tested for the existence of technological spillovers using a panel of U.K. manufacturing plants. Their analysis focused on whether the distance of domestic plants from the technological frontier affected the magnitude of technology transfer from high to less technology-intensive establishments. They found that technological spillovers from foreign multinationals serve as an important determinant of plant-level productivity convergence, as frontiers dominated by foreign plants bestowed larger productivity benefits on domestic producers.⁶

The principal aim of this paper is to test for establishment-level spillovers from growing to declining rivals. Baldwin (1995) found little evidence using industry-level average productivity for the two groups that productivity improvements in more successful producers filter down to less successful rivals. What follows is a more detailed plant-level analysis of the possible spillover effects associated with competitive reallocation.⁷ In this paper, we use the relationship between the productivity performance of declining and growing plants as evidence of these spillovers. Our analysis is founded on the notion that knowledge spillovers, if they are important, will be manifested in productivity performance. Our regressions evaluate whether the productivity performance of incumbents that are losing market share over three time periods is affected by (a) their economic and physical proximity to the technological frontier in the initial year of observation, and (b) contemporaneous productivity changes in rivals that are gaining market share over the period. The first of these possible senders of productivity spillovers, producers at the technological frontier, represents ‘established’ technology; the second, market-share gainers, represents ‘emerging’ technologies.

We are interested in these two possible sources of spillovers—frontier plants at the beginning of each period and plants gaining market share over the period—because of what they reveal about the types of information that firms draw on when attempting to improve their performance. The potential receivers of spillovers in our model are incumbent plants that have been ceding ground to competitors. These are establishments with a clear incentive to appropriate productive knowledge from others. To this end, highly efficient rivals located at the technological frontier at the beginning of a period represent a possible source of productive knowledge. Frontier plants amount to a visible set of mature industry leaders against which competitors can benchmark their relative performance. These are producers that may have established a more codifiable—and

6. Foreign firms are not the only source of technological spillovers. Griffith, Redding and Simpson (2002, p. 24) also report that “high productivity foreign-owned establishments are found to make an important contribution to productivity growth by pushing out the technological frontier, but high productivity UK-owned firms, to the extent that they are also technological leaders in some industries, are also important sources of technology transfer.” Baldwin and Gu (2005, p. 6) present similar findings, arguing that the “foreign ownership advantage is found to be a multinational advantage;” the authors note that Canadian multinationals operating in Canadian manufacturing industries exhibit technology, wage and productivity profiles on par with foreign multinationals.

7. We should note that testing for spillover effects was not the primary focus of Baldwin, as his analysis focused on estimating the contribution of plant turnover to aggregate productivity growth (see 1995, chap. 9).

appropriable—set of best practices. We examine whether a struggling plant’s initial proximity to its frontier affects its subsequent rate of productivity growth. Therefore, we are interested in the extent to which plants that are losing market share look to—or more accurately, look back to—the existing technological frontier for ideas on how to arrest their relative decline.

Market-share gainers are a second possible source of productivity spillovers: these are establishments that are actively wresting market share away from their rivals, thereby demonstrating a revealed advantage. These plants transmit a different type of productive knowledge to competitors—knowledge of emerging technologies and other business strategies that are currently associated with success and that are redefining the frontier. Here, productivity spillovers may depend on what struggling plants—the receivers of spillovers—can contemporaneously glean from their better-performing rivals. It is possible that this information is less appropriable than the best practices favoured by industry leaders because of its novelty. This said, plants that are actively wresting market share away from rivals would certainly provide struggling firms with a real-time source for productive learning.

Differentiating between productivity spillovers from the established frontier and from emerging market-share gainers affords us a clearer picture of the dynamic process at work as rivals continuously work to improve their competitive position. Less competitive producers—plants that are shedding market share—may look to established rivals at the technological frontier for productive information. These firms may also be able to use this information to improve their efficiency, either immediately or over time. But the set of technological best practices that form the basis for these spillovers is continuously being re-invented, either by existing members of the frontier and/or by innovative upstarts who push the boundaries of productive knowledge outward in new directions. Consequently, firms may be using older information from the frontier to catch up with established rivals, while at the same time falling behind other rivals that are actively breaking away from the pack by setting new competitive standards. Following Baldwin (1995), our distinction between market-share gainers—as potential senders of spillovers—and market-share decliners—as possible recipients of spillovers—is meant to draw attention to these emerging sources of knowledge embedded in growing firms and whether it is transferred to the group whose market-share performance suggests they might benefit most.

In what follows, we evaluate whether there is statistical evidence of productivity spillovers from more to less successful establishments. Several measurement issues warrant emphasis. First, it should be noted that the manufacturing panel used in this study does not provide observational data on the mechanisms by which productivity spillovers, if apparent, are being transmitted—such as through the learning and imitation of technological best practices, the dissemination of knowledge of exporting, or through pro-competition effects that encourage producers to improve their production efficiencies.⁸ This said, the concept of productivity spillovers advanced herein most closely aligns with the notions of active learning and imitation in which lesser performing establishments absorb technological knowledge from better-performing rivals. As with previous

8. Some analysis of transmission mechanisms is found in Baldwin and Gu (2005). Using a linked dataset that combines longitudinal data on plant performance with data collected from a technology survey, the authors test for the presence of spillovers by examining whether cross-industry variation in the share of output under foreign control influences (a) the level of competition faced by domestic plants, and (b) the extent to which these plants adopt advanced technologies. In both cases, the impact of foreign control is found to be positive and significant, suggesting that competition and technology use are areas in which domestic companies benefit from the co-existence of foreign-controlled plants.

studies, we rely on an observable metric—productivity growth—to infer something about the existence of these knowledge spillovers.

Second, we should stress that the productivity growth of declining plants, the metric on which our empirical analysis is based, can be influenced by different factors—some of which are consistent with the idea of spillovers from active learning and some of which may accrue more passively to these firms as a result of growth conditions within the industry. We are principally interested in the former—the localized productivity benefits that are transmitted to decliners from other segments of the plant population as proprietary knowledge from these producers spills over to rivals. These are spillovers in the strict sense. But the productivity performance of decliners is apt to be correlated with the underlying technological conditions that characterize the industry. Industry environments more conducive to growth may create ‘a rising tide’ for all members of the plant population, irrespective of the localized spillovers that occur within that population. These rising tides amount to a shared effect, common to gainers and decliners, that may have little to do with productive knowledge that spills over from some producers to others. We are mindful of this issue as we proceed, as our analysis is based on establishment-level data from over 200 four-digit manufacturing industries.

3 Data and measurement

Our data source is a representative panel of incumbent manufacturing plants derived from Statistics Canada’s Annual Survey of Manufactures, covering the period from 1979 to 1996. Constructed from annual surveys and administrative records, this panel—the Longitudinal Manufacturing Research File—has been used extensively to support micro-economic research on productivity.⁹ Our data file contains information on manufacturing production, value added, employment, inventories and four-digit industry of operation for all Canadian manufacturing plants in select years. It also includes longitudinal identifiers that allow for the identification of entrants, exits and incumbent plants. Our database also contains detailed spatial information on the location of competitor plants. These provide us with year-specific data on the distance of each individual plant in our sample to all other plants operating within a common four-digit industry.¹⁰

In what follows, we examine cross-sectional samples of incumbent manufacturing plants in three separate analysis periods: 1979 to 1984, 1984 to 1990 and 1990 to 1996. All units in these samples were plants with positive employment that operated in the same four-digit industry at the start and end years of the analysis period. Our choice of five- to six-year analysis intervals is designed to capture non-transitory changes in incumbent performance. Three intervals are examined separately to evaluate the sensitivity of our results to the choice of analysis period. Two intervals—1979 to 1984 and 1990 to 1996—correspond to periods with large increases in manufacturing productivity—as annualized labour productivity grew by 3.6% from 1979 to 1984

9. The version of the Longitudinal Manufacturing Research File (LMRF) used here covers the years from 1973 to 1999 inclusive with a constant Standard Industrial Classification industry code. Many of the studies that have utilized the LMRF are published in Statistics Canada’s Economic Analysis Research Paper Series.

10. These spatial data were developed at Statistics Canada to support a series of research projects that have examined the link between urban agglomeration and productivity growth (e.g., Baldwin et al. 2007; Baldwin, Brown and Rigby 2008).

and by 2.8% from 1990 to 1996. The middle period, covering the years from 1984 to 1990, saw weaker growth in manufacturing productivity, at 0.8% per annum.¹¹

Table 1
Cumulative shifts in market share, continuing plants

	1979 to 1984	1984 to 1990	1990 to 1996
		number	
Continuing plants	18,492	19,541	22,666
		percent	
Continuing plants that gain market share during analysis period			
Market share in first year of analysis period	36.8	33.6	36.7
Market share in last year of analysis period	49.7	44.7	53.4
Continuing plants that lose market share during analysis period			
Market share in first year of analysis period	49.4	52.6	44.1
Market share in last year of analysis period	36.8	38.9	32.8

Notes: All market shares are calculated in current dollars. Note that the combined market shares of gainers and decliners within a given analysis period do not sum to 100 because the total market share includes the shares held by non-continuing plants (entrants and exits).

Source: Statistics Canada, Longitudinal Manufacturing Research File.

We divide the continuing plant populations in each sample period into two groups: (1) plants that gain market share and (2) plants that lose market share over the interval.¹² A substantial amount of total market share changes hands as a result of these competitive transfers (Table 1). Incumbent plants that gained share over the 1979-to-1984 period began in 1979 with a cumulative market share of 36.8%; by 1984, these plants accounted for one half (49.7%) of all manufacturing production. Plants that lost share over this period, by contrast, saw their cumulative share of manufacturing production fall from 49.4% in 1979 to 36.8% in 1984. From 1990 to 1996, market-share gainers saw their cumulative share of manufacturing production increase by 17 percentage points—from 36.7% in 1990 to 53.4% in 1996. In contrast, market-share decliners saw their share of cumulative production decline from 44.1% to 32.8%.

The firms that gain market share do so as the result of considerable gains in relative productivity. At the beginning of an observation period, plants that go on to gain market share are not, on average, substantially more productive or less productive than plants that go on to lose share, but it is growing plants that, on average, increase their productivity dramatically over each period. Consider the sample of continuing plants operating from 1979 to 1984. In 1979, plants that were

11. See CANSIM table 383-0022. These estimates of annualized growth are based on value added.

12. For each plant i , this change is calculated as the first difference of i 's share of production within its four-digit industry of operation,

$$\Delta MS_{t_s, t_e}^i = \frac{VPM_{t_e}^i}{\sum_i^k VPM_{t_e}^i} - \frac{VPM_{t_s}^i}{\sum_i^k VPM_{t_s}^i}$$

where VPM is real manufacturing production, k is the number of plants in plant i 's industry, and t_e and t_s denote the end year and start year for the period. Note that these plant-specific shares for continuers are based on their portion of total manufacturing production, irrespective of plant types (i.e., $\sum_i^k VPM_{t_e}^i$ and $\sum_i^k VPM_{t_s}^i$ are based on all employer plants operating in either t_s or t_e , not just incumbent plants).

about to grow were 3% more productive than the incumbent average (the union of growing and declining plants); by 1984, these growers were 12% more productive.¹³ Similar shifts in relative productivity are apparent in the other analysis samples—with market-share gainers increasing their relative productivity by 8 percentage points in each period. Previous research has shown that the emergence of these differences in relative productivity is partly related to underlying differences in the capabilities of firms to innovate and exploit advanced technology (Baldwin and Sabourin 2001, 2004; Baldwin, Sabourin and Smith 2003; Baldwin and Gellatly 2003). Further, it is these changes in relative productivity that foster the observed changes in market share.

Table 2
Relative productivity of continuing plants

	1979 to 1984	1984 to 1990	1990 to 1996
	index ¹		
Continuing plants that gain market share during analysis period			
Labour productivity in first year of analysis period	1.03	0.97	0.95
Labour productivity in last year of analysis period	1.12	1.05	1.03
Continuing plants that lose market share during analysis period			
Labour productivity in first year of analysis period	0.98	1.02	1.05
Labour productivity in last year of analysis period	0.87	0.95	0.96

1. All continuing plants = 100.

Note: The productivity of continuing plants that gain or lose market share is calculated relatively to all continuing plants in the observation period.

Source: Statistics Canada, Longitudinal Manufacturing Research File.

While, on average, market-share gainers exhibit better productivity performance than market-share decliners, not all plants follow the same trajectory. Chart 1 reports on the productivity growth distributions of both groups, averaged across the three analysis periods.¹⁴ There is substantial plant-level variability in the productivity performance of both gainers and decliners. Plants that increased their market share exhibit, on balance, faster productivity growth than those with declining shares. On average, the median annual rate of productivity growth of plants gaining market share was about 4 percentage points higher than the median rate for plants losing share.¹⁵

13. For each of our analysis samples, the incumbent average is calculated as the weighted average level of labour productivity for all continuing plants operating during the observation period. This incumbent population is effectively the union of market-share gainers and market-share decliners as described herein, as these tabulations exclude the impact of entrants and exits.

14. Productivity growth for plant i , ΔLP^i , is calculated as the compound annual growth rate over the analysis period

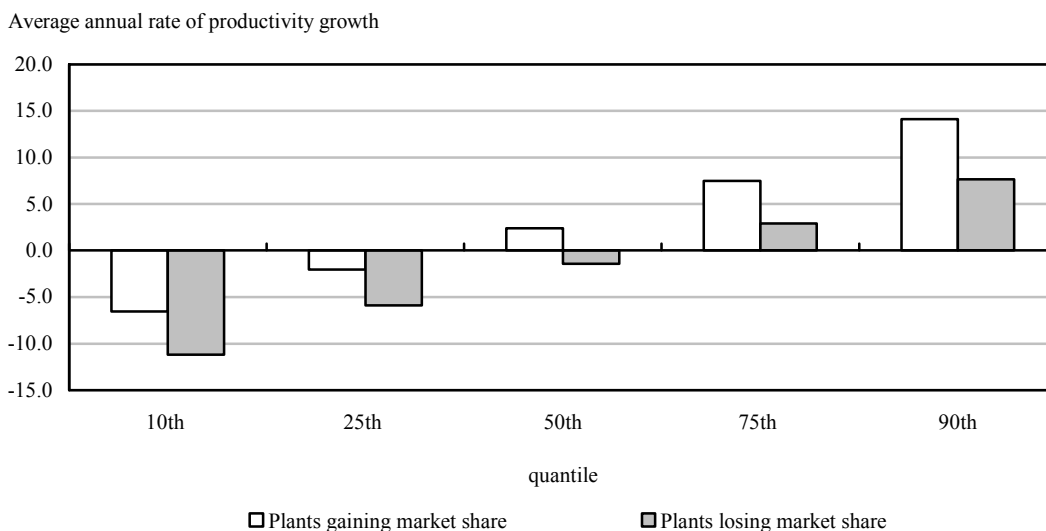
$$t_s, t_e: \Delta LP^i_{t_s, t_e} = \left[\left(\frac{LP^i_{t_e}}{LP^i_{t_s}} \right)^{\frac{1}{t_e - t_s}} \right] \times 100. \text{ Labour productivity for plant } i, LP^i, \text{ is the ratio of real manufacturing}$$

production VPM^i to plant employment EMP^i : $LP^i = \frac{VPM^i}{EMP^i}$. We use production per worker rather than

value added per worker in order to reduce the number of aberrant observations in the microdata due to reporting errors. Previous studies have produced quite similar results at a more aggregated industry level using both formulations and our experiments produced the same result.

15. Note that these descriptive tabulations do not take differences in industry membership into account.

Chart 1
Distribution of productivity growth rates by market share status
(all-period average)



Source: Statistics Canada, Longitudinal Manufacturing Research File.

Very sizable differences between gainers and decliners are also apparent at both ends of the productivity distribution—among plants in each group with high and low rates of productivity growth. Market-share gainers with faster productivity growth—those at the 75th and 90th percentiles of the growth rate distribution for gainers—enjoyed annual increases in labour productivity that were some two to three times larger than those of declining plants with high rates of productivity growth. Similarly, while market-share gainers at the low end of their productivity distribution—plants at the 25th and 10th percentiles—saw reductions in their labour productivity, these reductions were not nearly as severe as those apparent for declining plants at the low end of their productivity distribution.

These data afford us some insight into the dynamics of competitive restructuring at the plant level. At the beginning of each period, the productivity levels of gainers and decliners are quite similar on average. But over each study period, the productivity of market-share gainers forges ahead as the result of technological progress. Idiosyncratic productivity improvements in this part of the population that are passed on to consumers in the way of relative price or quality improvements create a centrifugal force that transfers market share from the less successful to the more successful.

Opposing this tendency are centripetal forces that help those losing market share to attenuate their productivity differences with those who are surging ahead. Externalities contribute to these centripetal forces. The special knowledge of techniques, technologies, patents and organizational forms that permits some firms to develop an advantage can be transmitted to those who fall behind. The magnitude and speed with which this transfer occurs will determine whether those who did not initially discover new techniques or new products survive in the market place.

In the next section, we examine the performance of the group of firms who are losing market share. These businesses, on average, have negative productivity growth. That their overall productivity performance is negative might be regarded as sufficient evidence that struggling

producers are unable to obtain productivity spillovers from more successful rivals. One could claim, on empirical grounds, that productive spillovers are probably absent from markets where declining plants—producers that find themselves on the wrong end of the competitive struggle—are performing poorly. If the majority of these producers are becoming less, not more, efficient as a result of competition, why would one expect that these declining plants are assimilating productive knowledge from rivals?

In our view, the productivity data on market-share gainers and decliners do not provide sufficient grounds for ruling out the existence of possible spillovers associated with competitive reallocation. First, while plants in decline have, on balance, experienced negative growth, their productivity performance could have been worse. Second, both groups—market-share gainers and market-share decliners—are heterogeneous populations, each encompassing a diverse range of productivity outcomes. It is not the case that all plants that are transferring market share to rivals are also becoming less productive, though many clearly are. Many struggling plants do succeed in improving their productive efficiency. In each of our samples, the declining plants located in the top quartile of their productivity-growth distribution all increased their labour productivity faster than the median market-share gainer.¹⁶ Thus, while many plants with declining output shares also became less efficient, others became more efficient. At issue is the extent to which the variation in productivity performance among these declining plants is consistent with an externality process wherein some struggling firms obtain productive knowledge from more successful rivals. This requires microdata on the performance of all decliners and an examination of how each was affected by the various potential sources of spillovers.

4 A model of productivity spillovers

In this section, we propose an estimation model that evaluates the extent to which the productivity growth of market-share decliners is influenced by (1) the economic and physical proximity of the declining plant to its technological frontier, and (2) contemporaneous productivity improvements in rival plants that have gained market share. Two basic assumptions underlie the model of the externality process advanced herein. First, spillovers flow unidirectionally from better performers—frontier plants and market-share gainers—to lesser performers—plants losing market share. Second, these spillovers are limited to plants operating in the same four-digit industry: cross-industry effects are not considered. To this spillover model, we also include a term that captures the basic regression-to-the-mean movements that have been previously found to be prevalent in growth processes.

Our estimation model can be represented as

$$\begin{aligned} \Delta lp_{t_s, t_e}^{ij} = & \alpha + \beta_1 \left[\overline{lp}_{t_s}^{f_j} - lp_{t_s}^{ij} \right] + \beta_2 \left[\left(\sum_{t_s}^{f_j} d_{t_s}^{ij \rightarrow f_j} \right) / n^{f_j} \right] \\ & + \beta_3 \left[\left(\sum_{ij \rightarrow g_j}^{g_j} (\Delta lp_{t_s, t_e}^{g_j} / d_{ij \rightarrow g_j}^{\delta}) \right) / n^{g_j} \right] + \beta_4 \left[lp_{t_s}^{ij} - \overline{lp}_{t_s}^{d_j} \right] + u \end{aligned} \quad (1)$$

16. Again, we reserve any rigorous consideration of industry effects for the formal econometric model.

where the left-hand-side term, $\Delta lp_{t_s, t_e}^{ij}$, is the annual growth in labour productivity of declining plant i in industry j over the observation period (t_s, t_e) .¹⁷

The first two right-hand-side terms evaluate the impact of the technological frontier on the productivity growth of declining plants. The first, $\left[\overline{lp}_{t_s}^{F_j} - lp_{t_s}^{ij} \right]$, measures the economic distance of the declining plant from its technological frontier F_j at the start of the period. This frontier F is common to all declining plants operating in four-digit industry j . To define F_j , all plants were ranked from highest to lowest, based on their level of labour productivity in the initial year of the analysis period, and were then classified sequentially to the frontier until the cumulative production of these plants accounted for at least 20% of the industry total.¹⁸ We then operationalized the frontier as the average level of labour productivity $\overline{lp}_{t_s}^{F_j}$ within the set of plants that define the frontier (see Baldwin 1992). A declining plant's economic distance from the frontier is thus defined as the difference between the productivity average for the frontier, $\overline{lp}_{t_s}^{F_j}$, and the productivity level of the declining plant, $lp_{t_s}^{ij}$.¹⁹ It is measured in constant dollar productivity units of \$1,000 per employee.

The coefficient on this frontier term, β_1 , captures the subsequent change in the productivity growth of the declining plant associated with the productivity gap between it and industry leaders (i.e., frontier plants) in the first year of observation. Plants that are located farther away from the frontier are hypothesized to grow more rapidly in subsequent years. This may occur if productivity spillovers accrue disproportionately to lower productivity producers. Our measure of economic distance is thus designed to measure a form of catching up—a narrowing of productivity differentials as technological knowledge from the frontier spills over to struggling rivals.

This measure of a declining plant's economic distance from the frontier $\left[\overline{lp}_{t_s}^{F_j} - lp_{t_s}^{ij} \right]$ assumes that productivity spillovers from frontier plants to decliners are not influenced by the latter's physical proximity to the frontier. This is, in our view, a sensible assumption if frontier plants constitute a well identified group of technology leaders that are readily observable by rivals. Frontier plants are the most productive producers in an industry, and higher levels of labour productivity in these plants should reflect observable differences in technology and capital intensity. Both local and distant rivals can be expected to monitor these industry leaders in an attempt to acquire productive information on new technologies and other best practices.

It is certainly possible, however, that the ability of struggling plants to incorporate productivity spillovers from the frontier is spatially circumscribed. We evaluate this via a separate variable that measures the average physical distance of each declining plant from frontier rivals. This

17. As noted in Section 3, our measure of productivity growth is the plant's compound annual rate of productivity growth over the observation period.

18. This ranking was based on all plants operating in the start year of the analysis period—both continuers and non-continuers.

19. The use of productivity gaps to operationalize the distance from the frontier is also employed by Griffith, Redding and Simpson (2002), who employ a measure of total factor productivity.

plant-specific measure is calculated as $\left[\left(\sum_{j \in \mathcal{F}} d_{t_s}^{i \rightarrow j} \right) / n^{\mathcal{F}} \right]$, where $\sum_{j \in \mathcal{F}} d_{t_s}^{i \rightarrow j}$ is the sum of all the individual distances from the declining plant i to all frontier plants j , divided by the number of frontier plants $n^{\mathcal{F}}$. This distance variable is measured in kilometres. If productivity spillovers from the frontier are constrained by distance, then β_2 should be negative.

Equation (1) treats the declining plant's economic distance from the frontier and its physical distance from the frontier as independent factors. Both effectively describe a backward-looking process in which the receivers of spillovers benefit from an earlier set of technological best practices established by frontier rivals at the start of an observation period t_s . These variables tell us little about the capacity of plants to assimilate new productive information as it becomes available.

The third right-hand-side variable in Equation (1) is designed to measure the spillovers associated with continuous learning from emerging technologies and organizational practices. It evaluates the productivity impact of market-share gainers on market-share decliners as industry output is re-allocated over the observation period $t_s \rightarrow t_e$. Here we impose stronger assumptions about how productive knowledge is being transmitted from more to less successful rivals. We posit that declining plants should seek to learn only from rival gainers that exhibit superior rates of productivity growth. Accordingly, the set of market-share gainers that can be expected to transmit spillovers to declining plants differs for each declining plant in our sample—with each decliner's rate of productivity growth acting as a threshold to determine which of its rival gainers are in-scope as potential senders. Declining plants with relatively high rates of productivity growth will have fewer rivals that can act as potential senders of spillovers, while decliners with low rates of productivity growth could draw from a larger set of potential senders.

We assume that differences in the causes of productivity growth are observable across rival plants. These productivity differentials between market-share gainers and market-share decliners reflect underlying differences in technology use, capital intensity and labour skill—characteristics that in many cases will be observable to firms as they benchmark their competitive progress against rivals. This said, we strongly expect that the ability of declining plants to acquire productive knowledge from market-share gainers is negatively constrained by distance. Gainers are not a small, well defined set of technology leaders, as is often true of frontier plants. Gainers are emerging plants that are wresting market share away from competitors—together these plants comprise about one half of the incumbent population and the causes of their sudden success may not immediately be perceived, comprehended or easily duplicated by their peers. We expect that declining plants that are looking to gainers for competitive knowledge are more apt to focus on and/or learn from competitors in local markets. These are rivals in close proximity that are actively improving their competitive position and that are likely to be far more observable to struggling plants than other rival gainers located in more distant geographic markets.

Our variable for measuring productivity spillovers from market-share gainers to market-share decliners is constructed as follows: for each declining plant, we calculate a distance-weighted average rate of productivity growth for rivals gaining market share. This is

$(\sum_{ij \rightarrow gj} (\Delta p_{t_s, t_e}^{gj} / d_{ij \rightarrow gj}^\delta)) / n^{gj}$ where $(\Delta p_{t_s, t_e}^{gj} / d_{ij \rightarrow gj}^\delta)$ is the annual rate of productivity growth of a rival market-share gainer g inversely weighted by its distance from the declining plant i . An average rate is generated by summing these plant-specific, distance-weighted growth rates for rival gainers and then dividing by the number of these rivals n^{gj} . We should again stress that the set of market-share rivals $g=1, \dots, G$ included in the above calculation will differ for each declining plant i . For each i , the candidate set of market-share gainers is defined as those that increased their productivity over the period $t_s \rightarrow t_e$ and that exhibited larger annualized rates of productivity growth than i ($\Delta p_{t_s, t_e}^{gj} > \Delta p_{t_s, t_e}^{ij}$).

The inverse distance weight $d_{ij \rightarrow gj}^\delta$ that is applied to the productivity-growth rate of each rival market-share gainer g assigns greater importance to gainers that are located closer to ij . In cases where the physical distance between ij and g is less than one kilometre, $d_{ij \rightarrow gj}^\delta$ is set to one. This effectively assumes that the productivity spillovers from gainers are fully appropriable by decliners located in very close proximity. In principle, the inverse distance weight also includes a decay parameter δ , which we effectively set to one. By so doing, we assume that the influence of a rival gainer gj on declining plant ij decays linearly as the distance between ij and gj increases.

This distance-weighted variable is designed to evaluate contemporaneous productivity spillovers that are associated with the transfer of market share from declining plants to gaining plants over the observation period $t_s \rightarrow t_e$. It yields an estimate of the amount of productivity growth generated by possible senders of spillovers, after taking into account the friction of distance (expressed as units of growth per kilometre). If declining plants are acquiring positive productivity spillovers from gainers, β_3 should be positive.

The fourth term in Equation (1), $[lp_{t_s}^{ij} - \bar{lp}_{t_s}^{d_j}]$, captures an important component of the stochastic process that underlies growth distributions—the tendency of firm-specific growth rates to regress to the mean. This has been shown to be a persistent characteristic of longitudinal data. These random movements do not conform with the concept of productivity spillovers in the sense that we are evaluating these herein. Plants that are generally ahead of the mean growth rate will tend to decline towards the mean over any period and vice versa. We evaluate this tendency towards mean reversion by controlling for the gap between the declining plant’s level of labour productivity in the initial period $[lp_{t_s}^{ij}]$ and the mean level of labour productivity exhibited by all declining plants within the declining plant’s industry $[\bar{lp}_{t_s}^{d_j}]$. In the presence of mean reversion, declining plants that start out being more productive than the average decliner (ostensibly the point around which random fluctuations occur) can be expected to experience slower growth rates than rival decliners that start out less productive than the average decliner. Because of these random movements, the coefficient on this term, β_4 , should be negative.

Our representation of the spillover process (Equation [1]) warrants some elaboration. It is purposely simplistic, in that it abstracts from plant-specific characteristics that may be correlated

with productivity differentials across producers, such as differences in capital intensity and/or labour quality. We have omitted these and other possible sources of organic productivity growth from our model, limiting its focus to correlative relationships among rival producers.

Equation (1), as presented above, also abstracts from industry fixed effects—more general factors that may be correlated with differences in productivity growth across manufacturing sectors. These factors can be envisaged as technological changes that benefit all producers in an industry, while conferring greater benefits on some industries than on others. They may correspond to basic differences in industry life cycle, with faster rates of technological change in sunrise industries fostering higher rates of productivity growth among both growing and declining plants. As these industry fixed effects are shared by growers and decliners, they seem far removed from the idea of competitive spillovers advanced herein, wherein some rivals, through learning, imitation or some other comparable process, extract productive knowledge from rivals.²⁰

There is evidence of these industry effects in each of our cross-sectional samples. Industries in which gainers exhibit relatively high rates of productivity growth tend also to be those in which decliners do comparatively well. In each analysis period, the average rates of productivity growth for market-share gainers and market-share decliners are positively correlated across industries. These industry averages yield correlation coefficients of 0.41 for the 1979-to-1984 period, 0.47 for the 1984-to-1990 period and 0.30 for the 1990-to-1996 period.

One strategy for controlling for these shared industry effects is to include industry-specific intercept terms in Equation (1). The estimation equation then becomes

$$\begin{aligned} \Delta p_{t_s, t_e}^{ij} = & \alpha + \sum \lambda_j + \beta_1 \left[\overline{lp}_{t_s}^{f_j} - lp_{t_s}^{ij} \right] + \beta_2 \left[\left(\sum_{i \rightarrow f_j} d_{t_s}^{i \rightarrow f_j} \right) / n^{f_j} \right] \\ & + \beta_3 \left[\left(\sum_{ij \rightarrow g_j} (\Delta p_{t_s, t_e}^{g_j} / d_{ij \rightarrow g_j}^{\delta}) \right) / n^{g_j} \right] + \beta_4 \left[lp_{t_s}^{ij} - \overline{lp}_{t_s}^{d_j} \right] + u \end{aligned} \quad (2)$$

where $\sum \lambda_j$ is a vector of four-digit industry controls, encompassing approximately 200 separate manufacturing industries. This specification allows the productivity growth of decliners to differ across four-digit industries, while assuming that the impact of productivity spillovers to declining plants from frontier rivals and market-share gainers—as measured by the slope terms $\beta_1, \beta_2, \beta_3$ —are equivalent from industry to industry. But this formulation is akin to demeaning each variable by its industry mean—which has direct implications for the separability of the frontier and mean reversion terms as expressed separately in Equation (2). When industry-specific intercepts are included in the model, the frontier and mean reversion terms collapse to a

20. We are tacitly assuming that the source of these shared effects is exogenous to both market-share gainers and market-share decliners. A full consideration of these shared effects—beyond our scope here—would necessarily ask whether they are truly exogenous. If some industries experience more rapid productivity growth than others, and this higher productivity emanates directly from the actions of market-share gainers, then these industry conditions do ostensibly represent a type of spillover—as gainers are creating the conditions that enable struggling firms to prosper. These distinctions underscore the challenges of conceptualizing spillovers within dynamic systems of competitive reallocation.

single regression-to-the-mean term $\left[lp_{t_s}^{ij} - \overline{lp}_{t_s}^{d_j}\right]$ —with a single coefficient that is equivalent to $\beta_4 - \beta_1$. Hence, this industry fixed effects formulation cannot be used to separately evaluate the impact of productivity spillovers from the frontier. For this reason, this fixed effect model is estimated without the frontier term.

5 Results and analysis

The model of the externality process described in Section 4 posits that the productivity growth of declining plants—those that lost market share over a five- to six-year observation period—depends separately on their economic distance and physical distance from their industry’s technological frontier—evaluated at the start of the observation period—and on contemporaneous productivity gains in plants that grew their market shares during the period.

Regression results for Equations (1) and (2) are reported in Table 3. The distribution of productivity growth rates for declining plants in each analysis period is close to being normal and therefore the ordinary least squares method is used for the estimation process. The first column for each analysis period contains the results of Equation (1) without the regression-to-the-mean term. The second column for each period contains the results of Equation (1) with the regression-to-the-mean term included. The third column reports the results for Equation (2), the industry fixed effects specification. This variant of our spillover model excludes the distance to the frontier term since it reduces to the difference from the mean term in this formulation. All p-values reported in Table 3 are based on robust standard errors, and in the case of the spatial covariates, clustered by geographic area²¹ to account for possible within-group correlations of error terms (see Moulton 1990).²²

Externalities, when they are significant, come generally from plants that are gaining market share. Plants with declining market shares obtain productivity spillovers from rivals gaining market share—the distance-weighted impact of productivity growth in market-share gainers (β_3) is positive and significant in the majority of cases. In the first two analysis periods, the impact of this growth-in-gainers term is significant in all formulations of our spillover model; in the last analysis period, these positive externalities from market-share gainers become significant, when industry fixed effects are taken into account. It is worth noting that, in all cases, these spillovers increase in magnitude once industry controls are included in the model (columns 3, 6 and 9).

There is little evidence from our estimation framework that a declining plant’s initial economic distance from frontier rivals has a positive impact on its subsequent rate of productivity growth. In the simplest formulation of our model, in which mean reversion and industry fixed effects are not taken into account, β_1 is positive and significant only in the last observation period (columns 1, 4 and 7). Once the tendency of these growth rates to regress to the mean is taken into account, β_1 becomes negative and significant in the first two observation periods, suggesting that less productive plants that start off behind the frontier go on to experience additional declines.

21. Geographic areas were defined as census metropolitan area/census area/census divisions.

22. A detailed discussion of the edit rules that were applied to our data samples, along with descriptive statistics, is found in Appendix A.

There is also little evidence that a declining plant's physical proximity to frontier rivals has an independent effect on its subsequent rate of productivity growth. A negative relationship between the decliner's growth rate and its average physical distance from frontier plants—one consistent with our hypothesis that $\beta_2 < 0$ —was apparent only in the final observation period (1990 to 1996). No significant relationship was observed in the first two periods.

Finally, the coefficient that evaluates the difference between the initial productivity of the declining plant and the mean productivity of this declining group is consistently negative and significant. This occurs irrespective of whether industry fixed effects are taken into account. The productivity growth of plants in decline follows a stochastic process with a strong reversion to the mean effect.

Table 3
Regression results, productivity growth of declining plants

	1979 to 1984			1984 to 1990			1990 to 1996		
	Column 1	Column 2	Column 3	Column 4	Column 5	Column 6	Column 7	Column 8	Column 9
Regressors									
Intercept									
Estimate (decimal notation)	-1.16	-0.26	...	-2.42	-1.44	...	-1.81	-1.29	...
P-value	0.02	0.59	...	0.00	0.00	...	0.00	0.01	...
Economic distance of the declining plant from its technological frontier in the initial period									
Estimate (scientific notation ¹)	-2.45E-04	-6.05E-03	...	2.19E-03	-4.16E-03	...	3.01E-03	4.15E-04	...
P-value	0.85	0.00	...	0.23	0.00	...	0.01	0.69	...
Average physical distance of the declining plant from frontier plants									
Estimate (scientific notation ¹)	-4.350E-05	-7.640E-05	-2.14E-04	1.94E-04	1.660E-04	-3.61E-05	-2.72E-04	-3.02E-04	-2.58E-04
P-value	0.83	0.71	0.19	0.29	0.34	0.78	0.01	0.01	0.03
Average productivity of gainers with higher productivity than the declining plant (inversely weighted by their distance to the decliner)									
Estimate (decimal notation)	0.63	0.63	1.51	1.34	1.40	1.96	0.73	0.81	1.73
P-value	0.01	0.01	0.00	0.01	0.01	0.00	0.24	0.18	0.00
Regression to the mean of decliners									
Estimate (scientific notation ¹)	...	-3.32E-02	-2.71E-02	...	-3.61E-02	-3.22E-02	...	2.27E-02	-2.32E-02
P-value	...	0.00	0.00	...	0.00	0.00	...	0.00	0.00
Diagnostic statistics									
Industry controls	None	None	4-digit	None	None	4-digit	None	None	4-digit
Observations (number)	8,174	8,174	8,174	9,066	9,066	9,066	10,279	10,279	10,279
R ²	0.0042	0.0599	0.1638	0.0134	0.0990	0.1784	0.0135	0.0583	0.1532

1. The figures are followed by the letter E, which stands for 'exponent', a plus sign or a minus sign, and two digits that indicate the power of 10 by which the number is multiplied. This is a way to write numbers that accommodate values too large or too small to be conveniently written in decimal notation.

Notes: The dependent variable is *average annual rate of labour productivity growth, plants with declining market shares*. P-values reflect standard errors that are corrected for heteroskedasticity and, in the case of spatial covariates, clustered by geography (census metropolitan area/census area/census division). Standard errors on our measure of economic distance are clustered by four-digit industry.

Source: Statistics Canada, Longitudinal Manufacturing Research File.

The size of the externalities in our model warrants emphasis. While there is evidence that plants in decline obtain productivity spillovers from rivals acquiring market share, their magnitude is relatively small. The mean productivity impact of gainers on decliners is presented in Table 4, along with the mean rate of productivity growth for decliners in each period. The latter is negative in all cases, and averages -1.4% per year over the three observation periods. The impact of the positive externality received from plants that are improving their productivity is relatively small—about 0.18 percentage point per year, on average, when industry fixed effects are not taken into account, and 0.34 percentage point per year, on average, when industry controls are included in our model. These productivity spillovers are relatively small when viewed in relation to the mean annual change in the productivity of declining plants (-1.4%). This suggests that without the impact of the externality from gainers, the rate of decline would have been -1.58% or -1.74% per year on average for each of these two estimates.²³

While the externality arising from the growing group is positive, it does not have much impact on the productivity differential between growing and declining plants. After the externality is fully internalized by decliners, there is still about a 3.5-percentage-point difference in the median growth rates of growing and declining plants over the period—a difference that would have only been slightly larger (at most 10%) if the externalities had not been operating. Externalities then have a relatively minor impact on closing the gap between more and less successful plants.

Table 4
Average impact of spillover from rivals gaining market share

Variable	1979 to 1984	1984 to 1990	1990 to 1996	Three-period average
	percent			
Average annual rate of productivity growth of the declining plant	-1.10	-1.60	-1.50	-1.40
	percentage points			
Average annual productivity growth arising from rival market-share gainers (no industry controls)	0.16	0.26	0.12	0.18
Average annual productivity growth arising from rival market-share gainers (with industry controls)	0.38	0.36	0.26	0.34

Source: Statistics Canada, authors' calculations.

We also examine the relative importance of each of the externalities by dividing the sample of four-digit industries into five major groups based on Equation (1). There is the possibility that part of the gainers effect observed in the aggregate sample is due to the fact that all plants are doing better in some industries than others. Accordingly, we divide manufacturing industries into five subsectors. These are natural resource-based industries, labour-intensive industries, scale-based industries, product-differentiated industries and science-based industries.²⁴ This classification system groups industries into different strata based largely on differences in the factors that influence the process of competition. For resource-based industries, the primary determinant of competition is access to abundant natural resources. For the labour-intensive sector, it is labour costs. For scale-based industries, competition hinges on the length of

23. We have ignored the impact of the frontier because it has a large negative effect when it is significant, thereby suggesting that it subtracts from rather than contributes to the centripetal effect of the externality.

24. See Baldwin and Rafiquzzaman (1994) for a discussion of the methodology used to create these groupings.

production runs. In the product-differentiated group, competition depends on an ability to target production to the demands of various markets. Competition in science-based sectors depends on the application of scientific knowledge. The latter two industries rely heavily on intangible brand and knowledge assets.

As was the case for the regressions for all sectors taken together, the frontier effects were generally not significant so they have been ignored here. But the coefficients that evaluate the externalities from market-share gainers were significant in 13 of the 15 period/industry cells. The three-period average for each of the industry groups is presented in Table 5 for the estimates without industry controls. As might be expected, the impact of the externality from market-share gainers differs across sectors. It is highest in the product differentiated and the science-based sectors, where knowledge of brands and technology are perhaps most important. It is lowest in the labour-intensive sector, where this type of imbedded knowledge is least important.

Table 5
Average impact of spillover from rivals gaining market share (three-period average)

Variable	Subsectors of manufacturing industries				
	Natural resource-based industries	Labour-intensive industries	Scale-based industries	Product-differentiated industries	Science-based industries
			percent		
Average annual rate of productivity growth of the declining plant	-1.52	-1.44	-1.30	-1.42	-1.21
			percentage points		
Mean impact of productivity growth of rival market-share gainers	0.39	0.14	0.47	0.83	0.85
			percent		
Average annual productivity growth of decliners without productivity spillovers	-1.91	-1.58	-1.77	-2.25	-2.06

Note: These results are derived from Equation (1) without the regression-to-the-mean term.
Source: Statistics Canada, authors' calculations.

The estimates provided in Table 4 summarize the impact of externalities on the 'average' declining plant. Alternately, we can evaluate the extent to which these externalities affect the productivity of declining plants by testing if large qualitative shifts in the exposure of declining plants to rivals gaining market share (our proxy for these productivity spillovers) has much of an impact on the relative productivity growth of decliners.

One means of investigating the relative size of these productivity impacts is to estimate the implied change in a declining plant's productivity associated with shifts across the interquartile range of the variable representing the productivity growth of market-share gainers—movements from the 25th percentile to the 75th percentile of its distribution. These represent qualitatively large changes in the amount of knowledge associated with rivals gaining share—potential senders of spillovers—that is available to a declining plant—as potential receivers of spillovers. Declining plants that are relatively well positioned at the 75th percentile face a collection of rival gainers that, on average, exhibit a more desirable combination of growth and distance characteristics. Those located at the 25th percentile compete against rivals with a less desirable combination of growth and distance characteristics.

Here some explanation is warranted. The idea of desirability depends on the amount of productivity growth that is generated by rival gainers that, in turn, can be transmitted to declining plants after accounting for the friction of distance. It is decliner-specific. A hypothetical decliner that is relatively well-positioned at the 75th percentile may be one with rivals, located nearby, that exhibit only modest rates of productivity growth. While these rivals may generate fewer spillovers than faster growing rivals, none of this productive knowledge is lost to the decliner plant because of its close proximity to these rivals. The decliner is well-positioned because its location facilitates the acquisition of this technological knowledge. An equally well-positioned decliner may compete against rivals with faster rates of productivity growth, but that are located much further away. These rivals may have larger spillovers to transmit; however, some of this productive knowledge will be lost (to the declining plant) over distance. It is the combination of growth and distance that determines the amount of productivity growth that is transmitted to decliners. Decliners located at the other end of the interquartile range, at the 25th percentile, have a substantially lower exposure to productivity spillovers from rival gainers because of a less advantageous combination of rival growth and distance characteristics. Shifting from the 25th to the 75th percentile represents a qualitative change in a declining plant's exposure to productivity spillovers.

We estimated the change in decliner productivity associated with these interquartile movements, based on the regression coefficients used to generate the average effects reported in Table 4. The exposure of declining plants to productivity improvements in rivals with increasing market shares has a small effect on decliner productivity. Shifts from relatively low exposure (the 25th percentile) to higher exposure (the 75th percentile) yield modest improvements in decliner productivity—0.13 percentage point per year on average, in the absence of industry controls, and 0.25 percentage point per year on average, when industry controls are taken into account.

6 Conclusions

Comprehensive microdata on business populations have led to highly detailed studies of competitive restructuring. These data sources reveal that competition is a dynamic process characterized by complementary forces of growth and decline. Industries undergo a continuous process of competitive replacement wherein new firms challenge and supplant established producers. In this paper, we have used panel data on Canadian manufacturing plants to study one aspect of this dynamic process—the productivity spillovers that may be occurring at the same time as the competitive reallocation of market share among rival incumbents.

The data examined in this paper show that even among continuing plants, large amounts of total market share are transferred among rival producers over relatively short periods of time. In each of the five- to six-year intervals studied here, producers that gained market share during the period finished with much larger cumulative shares of industry output with which they had begun, while decliners bore substantial losses. Plants on the right side of the competitive struggle often exhibit much stronger productivity growth than do those in decline. Standard productivity decomposition analyses show that this group accounts for the majority of productivity growth for two reasons: first, their organic productivity growth is largest; second, their higher organic growth results in their gaining market share. This shift of market share from less productive to more productive plants also creates a substantial amount of additional productivity growth.

This process creates centrifugal forces that separate the more productive from the less productive, increasing the market share for the former and decreasing the market share for the latter. Opposing this tendency are the centripetal forces that attenuate the growth of firms who initially surge ahead as a result of introducing new ideas and innovations that lead to higher productivity.

Externalities contribute to centripetal forces. The special knowledge of techniques, technologies, patents and organizational forms that permit some firms to develop an advantage can be transmitted to those which fall behind. The magnitude and speed with which this transfer occurs will determine whether those who did not initially discover new techniques or new products remain competitive.

The relative size of the centrifugal and centripetal forces will determine the nature of turnover—whether the forces that lead some firms to innovate and improve their productivity and market share are disseminated relatively quickly to those who are less agile in the innovative process initially. It will determine whether the Darwinian process of sorting resembles an ecological model where the firms with inferior productivity are sorted out and culled, or whether there is an innate support mechanism that pulls them along in order to keep up with the industry average.

To explore the existence of externalities that temper the productivity differences that develop between growing and declining plants, the paper has examined whether the competitive reallocation of market share involves productivity spillovers from more successful to less successful producers. It examines whether plants in decline obtain productivity spillovers more from old knowledge that is associated with the established frontier or from what is revealed by real-time performance to be successful new knowledge, which is at the edge of new, cutting-edge technology (Baldwin and Sabourin 2001). The latter is represented here by the productivity growth of rivals which are actively wresting market share away from decliners—plants whose growth provides evidence of a real advantage. These rivals are plants on the right side of the competitive struggle as the dynamics of competition are actively being played out.

Our results suggest that plants in decline obtain some measure of benefit from rivals, but they come mainly from spillovers emanating from market-share gainers to market-share decliners; however, the amount of externalities are spatially circumscribed and qualitatively modest, thereby suggesting that market-share gainers owe their advantages to knowledge that is more embedded and is less likely to be transferred freely in the form of externalities.

Plants losing market share do not obtain productivity spillovers from the old knowledge of rivals located at the technological frontier. These frontier rivals are high efficiency producers, as identified at the onset of each observation period. These are rivals which have previously established themselves as industry leaders. We find that declining plants starting out further behind the frontier actually do worse, not better. Declining plants that are further behind the frontier are more likely to decline than grow. Falling behind the frontier is not a winning strategy that firms can rely on to help them catch up in the next period.

The most important finding is that the amount of spillovers emanating from the plants experiencing productivity growth to those in decline is relatively minor. At most, the spillovers measured here increase the rate of productivity growth by 0.18 to 0.34 of a percentage point. This is relatively small compared to the nearly 4-percentage-point difference between the median growth rates of growing and declining plants over the period. The externalities measured here

then have a relatively minor impact on closing the gap between the more and the less successful plants.

These results then suggest a world where a group of firms find new ways to increase their productivity (perhaps serendipitously) and forge ahead of those with less luck or ability. While some of the newfound knowledge spills over the less capable or less lucky, the trickle-down effect is not sufficient to reduce the productivity differences that emerge. The evolutionary process resembles a contest that quickly rewards the more successful and leaves the less successful to fall further and further behind the leaders.

Appendix A Editing data samples for the regression analysis

An exploratory analysis was conducted to ascertain the extent to which the results of the externality model reported in Section 4 are sensitive to outliers. It suggested that outlier observations can have a qualitative impact on the results. Accordingly, outlier observations were removed from the data samples that were used for the regression analysis. These sample adjustments are described below, followed by descriptive statistics on the edited regression variables used to estimate Equations (1) and (2).

The spillover model advanced herein posits that the productivity growth of a plant with declining market share depends on its economic distance from the technological frontier, on its physical proximity to frontier rivals, and on productivity gains in rival plants gaining market share after accounting for the friction of distance. All of these factors are plant-specific.

Outlier observations were removed from several of these analysis variables in each analysis period. In the case of the dependent variable $\Delta p_{t_s, t_e}^i$, the tails of its distribution were trimmed at the 2nd and 98th percentiles. The variable measuring the declining plant's economic distance from the frontier $\left[\bar{l}p_{t_s}^F - lp_{t_s}^i \right]$ was trimmed at the 98th percentile, while the lower tail was truncated to eliminate small numbers of observations in each sample where the calculated distance from the frontier was non-positive—situations in which the declining plant started off with higher productivity than the average frontier plant. Edits on the productivity growth of rival plants gaining market share were conducted at the industry level. For each industry, we calculated the average rate of productivity growth for plants gaining market share from t_s to t_e , and then used these estimates to trim our analysis samples at the 2nd and the 98th percentiles—effectively eliminating declining plants located in industries in which gainers exhibited excessively high or low rates of productivity growth.

The remaining observations in each period were used to estimate the regressions reported in Table 3. These regression samples amounted to 87% of all observations on declining plants during the 1979-to-1984 period, and 88% of all observations on declining plants during the 1984-to-1990 and 1990-to-1996 periods, respectively. Descriptive statistics are reported in Table A.1.

Table A.1
Selected summary statistics, regression samples

	1979 to 1984	1984 to 1990	1990 to 1996
Variables			
Average annual rate of productivity growth of the declining plant ¹			
Mean value (percentage)	-1.1	-1.6	-1.5
Standard deviation	7.4	6.1	6.1
Economic distance of the declining plant from the technological frontier ²			
Mean value (constant-dollar productivity units of \$1,000 per employee)	148.4	149.4	189.6
Standard deviation	139.3	130.9	180.5
Average distance of the declining plant from frontier rivals ³			
Mean value (kilometres)	1,319.5	1,288.0	1,305.2
Standard deviation	813.6	746.7	809.4
Productivity growth of rival market-share gainers ⁴			
Mean value (percentage growth rate per kilometre)	0.25	0.19	0.15
Standard deviation	0.75	0.51	0.50
Diagnostic statistic			
Observations (number)	8,174	9,066	10,279

1. A compound annual growth rate calculated over the analysis period.

2. The average level of labour productivity of all plants located in the frontier (of the declining plant's industry) minus the labour productivity of the declining plant, both evaluated at the start of the observation period.

3. The cumulative distance to frontier rivals, divided by the number of frontier rivals.

4. An average measure of the distance-adjusted rate of productivity growth of plants gaining market share in the declining plant's industry. This average is based on the subset of gainers that have positive rates of productivity growth that exceed the productivity growth rate of the declining plant.

Note: The percentages of observations on declining plants included in the regression sample are as follows: for 1979 to 1984: 86.8%; for 1984 to 1990: 87.9%; for 1990 to 1996: 88.2%.

Source: Statistics Canada, Longitudinal Manufacturing Research File.

References

- Baldwin, John R. 1992. "Industrial efficiency and plant turnover in the Canadian manufacturing sector." In *Industrial Efficiency in Six Nations*. Richard E. Caves (ed.). Cambridge, Mass.: The MIT Press.
- Baldwin, John R. 1995. *The Dynamics of Industrial Competition: A North American Perspective*. Cambridge, U.K.: Cambridge University Press.
- Baldwin, John R., Desmond Beckstead, W. Mark Brown and David L. Rigby. 2007. *Urban Economies and Productivity*. Economic Analysis (EA) Research Paper Series. Catalogue no.11F0027MIE2007045. Ottawa: Statistics Canada.
- Baldwin, John R., W. Mark Brown and David L. Rigby. 2008. *Agglomeration Economies: Microdata Panel Estimates from Canadian Manufacturing*. Economic Analysis (EA) Research Paper Series. Catalogue no. 11F0027MIE2008049. Ottawa: Statistics Canada.
- Baldwin, John R., and Naginder Dhaliwal. 2001. "Heterogeneity in labour productivity growth in manufacturing: Differences between domestic and foreign-controlled establishments." In *Productivity Growth in Canada*. 61–76. Catalogue no. 15-204-XIE1999000. Ottawa: Statistics Canada.
- Baldwin, John R., and Guy Gellatly. 2003. *Innovation Strategies and Performance in Small Firms*. Northampton, Mass.: Edward Elgar Publishing.
- Baldwin, John R., and Guy Gellatly. 2007. *Global Links: Multinationals in Canada: An Overview of Research at Statistics Canada*. The Canadian Economy in Transition Research Paper Series. Catalogue no. 11-622-MIE2007014. Ottawa: Statistics Canada.
- Baldwin, John R., and Wulong Gu. 2005. *Global Links: Multinationals, Foreign Ownership and Productivity Growth in Canadian Manufacturing*. The Canadian Economy in Transition Research Paper Series. Catalogue no. 11-622-MIE2005009. Ottawa: Statistics Canada.
- Baldwin, John R., and Wulong Gu. 2006. "Plant turnover and productivity growth in Canadian manufacturing." *Industrial and Corporate Change*. 15, 3: 417–465.
- Baldwin, John R., and Wulong Gu. 2008. *Firm Turnover and Productivity Growth in the Canadian Retail Trade Sector*. Economic Analysis (EA) Research Paper Series. Catalogue no. 11F0027MIE2008053. Ottawa: Statistics Canada.
- Baldwin, John R., and Petr Hanel. 2003. *Innovation and Knowledge Creation in an Open Economy: Canadian Industry and International Implications*. Cambridge, U.K.: Cambridge University Press.
- Baldwin, John R., and M. Rafiquzzaman. 1994. *Structural Change in the Canadian Manufacturing Sector, 1970-1990*. Analytical Studies Branch Research Paper Series. Catalogue no.11F0019M, no. 61. Ottawa: Statistics Canada.

Baldwin, John R., and David Sabourin. 2001. *Impact of the Adoption of Advanced Information and Communication Technologies on Firm Performance in the Canadian Manufacturing Sector*. Analytical Studies Branch Research Paper Series. Catalogue no. 11F0019MIE, no. 174. Ottawa: Statistics Canada.

Baldwin, John R., David Sabourin and David Smith. 2003. *Impact of Advanced Technology Use on Firm Performance in the Canadian Food Processing Sector*. Economic Analysis (EA) Research Paper Series. Catalogue no. 11F0027MIE2003012. Ottawa: Statistics Canada.

Baldwin, John R., and David Sabourin. 2004. *The Effect of Changing Technology Use on Plant Performance in the Canadian Manufacturing Sector*. Economic Analysis (EA) Research Paper Series. Catalogue no. 11F0027MIE2004020. Ottawa: Statistics Canada.

Bartelsman, Eric J., and Mark Doms. 2000. "Understanding productivity: Lessons from longitudinal microdata." *Journal of Economic Literature*. 38, 3: 569–594.

Caves, Richard E. 1974. "Multinational firms, competition and productivity in host-country markets." *Economica*. 41, 162: 176–93.

Doms, Mark E., and J. Bradford Jensen. 1998. "Comparing wages, skills, and productivity between domestically and foreign-owned manufacturing establishments in the United States." In *Geography and Ownership as Bases for Economic Accounting*. Robert E. Baldwin, Robert E. Lipsey and J. David Richardson (eds.). Chicago: University of Chicago Press.

Ericson, Richard, and Ariel Pakes. 1995. "Markov-perfect industry dynamics: A framework for empirical work." *Review of Economic Studies*. 62, 210: 53–82.

Foster, Lucia, John Haltiwanger and C.J. Krizan. 2001. "Aggregate productivity growth: Lessons from microeconomic evidence." In *New Developments in Productivity Analysis*. Charles R. Hulten, Edwin R. Dean and Michael J. Harper (eds.). Chicago: University of Chicago Press.

Globerman, Steven. 1979. "Foreign direct investment and 'spillover' efficiency benefits in Canadian manufacturing industries." *Canadian Journal of Economics*. 12, 1: 42–56.

Globerman, Steven, John C. Ries and Ilan Vertinsky. 1994. "The economic performance of foreign affiliates in Canada." *Canadian Journal of Economics*. 27, 1: 143–157.

Griffith, Rachel, Stephen J. Redding and Helen Simpson. 2002. *Productivity Convergence and Foreign Ownership at the Establishment Level*. CEPR Discussion Paper no. 3765. London, U.K.: Centre for Economic Policy Research

Jaffe, Adam B., Manuel Trajtenberg and Rebecca Henderson. 1993. "Geographic localization of knowledge spillovers as evidenced by patent citations." *The Quarterly Journal of Economics*. 108, 3: 577–598.

Jovanovic, Boyan. 1982. "Selection and the evolution of industry." *Econometrica*. 50, 3: 649–670.

Markusen, James R. 1998. *Multilateral Rules on Foreign Direct Investment: The Developing Countries' Stake*. Mimeo. Washington, D.C.: World Bank.

Moulton, Brent R. 1990. "An illustration of a pitfall in estimating the effects of aggregate variables on micro units." *The Review of Economics and Statistics*. 72, 2: 334–338.

Organisation for Economic Co-operation and Development (OECD). 2001. "Productivity and firm dynamics: Evidence from microdata." *OECD Economic Outlook*. 69, 7. Paris: OECD.